



HAL
open science

Accelerating Iterative SpMV for Discrete Logarithm Problem using GPUs

Hamza Jeljeli

► **To cite this version:**

Hamza Jeljeli. Accelerating Iterative SpMV for Discrete Logarithm Problem using GPUs. 2012. hal-00734975v1

HAL Id: hal-00734975

<https://inria.hal.science/hal-00734975v1>

Preprint submitted on 25 Sep 2012 (v1), last revised 4 Dec 2014 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Accelerating Iterative SpMV for Discrete Logarithm Problem using GPUs

Abstract—In the cryptanalytic context, computing discrete logarithms in large cyclic groups using index-calculus-based methods, such as the number field sieve or the function field sieve, requires solving large sparse systems of linear equations modulo the group order. Most of the fast algorithms used to solve such systems — e.g., the conjugate gradient or the Lanczos and Wiedemann algorithms — iterate a product of the corresponding sparse matrix with a vector (SpMV). This central operation can be accelerated on GPUs using specific computing models and addressing patterns, which increase the arithmetic intensity while reducing irregular memory accesses. In this work, we investigate the implementation of SpMV kernels on NVIDIA GPUs, for several representations of the sparse matrix in memory. We explore the use of Residue Number System (RNS) arithmetic to accelerate modular operations. We target linear systems arising when attacking the discrete logarithm problem on groups of size 160 to 320 bits, which correspond to the sizes in the last record computations. The proposed SpMV implementation on a 1.7M-by-1.7M matrix containing 86M non-zeros, delivers a throughput of 24.3 SpMV/s on an NVIDIA GeForce GTX 580, which corresponds to 50 GFLOP/s.

Index Terms—Discrete Logarithm Problem, Sparse-Matrix-Vector product, Modular Arithmetic, Residue Number System, GPUs

I. INTRODUCTION

The security of many cryptographic protocols used for authentication, key exchange, encryption, or signature, depends on the difficulty of solving the discrete logarithm problem (DLP) in a given cyclic group [1]. For instance, we can rely on the hardness of the DLP in a multiplicative subgroup of a finite field. There are algorithms, such as Pollard-rho [2], Baby-Step/Giant-Step [3] that solve the problem in time exponential in the subgroup size. Another family of methods, known as *Index-calculus* [4] methods propose to solve it in time sub-exponential in the finite field size. These algorithms require in their linear algebra step the resolution of large sparse systems of linear equations modulo the group order [5]. In cryptographic applications, the group order ℓ is of size 160 to 320 bits. The number of rows and columns of the corresponding matrices are in the order of hundreds of thousands to millions, with only hundreds or fewer non-zero elements per row. This step is a serious limiting factor for the resolution. For example, it was reported in [6] that the linear algebra step of the Function Field Sieve (FFS) implementation to solve the DLP over $\text{GF}(3^{6 \times 97})$ took 80.1 days on 252 CPU cores, which represents 54% of the total time.

To solve such systems, ordinary Gaussian elimination is inefficient. While some elimination strategies aiming at keeping the matrix as sparse as possible can be used to reduce the input system somewhat, actual solving calls for the use of other techniques (Lanczos [7], Wiedemann [8]) that take advantage of the sparsity of the matrix [9]. Either for Lanczos,

Wiedemann or their block variants, the iterative sparse-matrix-vector product is the most time-consuming operation. For this reason, we investigate accelerating this operation on GPUs.

The paper is organized as follows. Section II presents the background related to the hardware and the context. Section III discusses the arithmetic aspects of our implementation. We present several matrix formats and their corresponding implementations in Section IV. We compare the results in Section V, and present optimizations based on hardware considerations in Section VI. Finally, Section VII concludes the paper.

II. BACKGROUND

A. GPUs and the CUDA programming model

CUDA is the hardware and software architecture that enables NVIDIA GPUs to execute programs written with C, C++, OpenCL and other languages [10].

A CUDA program instantiates a *host* code running on CPU and a *kernel* code running on GPU. The kernel code runs according to the Single Program Multiple Threads (SPMT) execution model across a set of parallel threads. The threads are executed in groups of 32, called *warps*. As a warp only has a single instruction fetch/decode unit, each instruction in the execution path is issued to all the threads in the warp. However, if one or more threads have a different execution path, execution divergence occurs. The different paths will then be serialized, negatively impacting the performance.

The threads are further organized into thread *blocks* and *grids* of thread blocks:

- A thread executes an instance of the kernel. It has a unique thread ID within its thread block, along with registers and private memory.
- A thread block is a set of concurrent threads which can share data through *shared memory* and perform barrier synchronization which ensures that all threads within that block fully reach the same instruction before continuing. It has a unique block ID within its grid.
- A grid is an array of thread blocks executing the same kernel. All the threads of the grid can also read inputs, and write results to *global memory*.

At the hardware level, the blocks are distributed on an array of multi-core *Streaming Multiprocessors* (SMs). Each SM schedules and launches the threads in groups of warps. Recent GPUs (such as the NVIDIA Fermi architecture) allow for up to 48 active warps per SM. The ratio of active warps to the maximum supported is called *occupancy*. Maximizing the occupancy is important, as it helps hiding the memory latency. One should therefore pay attention to the usage of shared memory and registers in order to maximize occupancy.

Another important performance consideration in programming for the CUDA architecture is *coalescing* global memory accesses. To understand this requirement, global memory should be viewed in terms of aligned segments of 32 words of 32 bits each. Memory requests are serviced for one warp at a time. If the warp request hits exactly one segment, the access is *fully coalesced* and there will be only one memory transaction performed. If the warp accesses scattered locations, the accesses are *uncoalesced* and there will be as many transactions as the number of hit segments. Consequently, the kernel should use a coalescing-friendly pattern for greater memory efficiency.

Despite their high arithmetic intensity and their large memory bandwidth, GPUs provide small caches. In fact, Fermi GPUs provide the following levels of cache for the on-board DRAM (local and global memory):

- 768-kByte *L2-cache* per GPU.
- 16-kByte *L1-cache* (per SM). It can be extended to 48 kBytes, but this decreases shared memory from 48-kByte to 16-kByte.
- A *texture cache*: an on-chip cache for the read-only *texture memory*. It can accelerate the memory accesses when neighboring threads read from nearby addresses.

B. Sparse-Matrix–Vector product on GPUs

Sparse-matrix computations pose some difficulties on GPUs, such as irregular memory accesses, load unbalance and low cache efficiency. Many works have focused on choosing suitable matrix formats and appropriate kernels to overcome the irregularity of the sparse matrix [11], [12], [13]. These works have explored implementing efficiently SpMV for numerical values. Schmidt et al. [14] proposed an optimized matrix format to accelerate exact SpMV over GF(2), that can be used in the linear algebra step of the Number Field Sieve (NFS) for integer factorization. Boyer et al. [15] have adapted SpMV kernels over small finite fields and rings $\mathbb{Z}/m\mathbb{Z}$, where they used a double to represent the ring. In our context, since the order of the finite ring is large (hundreds of bits), specific computing models and addressing models should be used.

In this work, we consider the product $w \leftarrow A \times v$, where A is a N -by- M sparse matrix. A feature of the application context is that A contains small values (e.g. 32-bit integers). v and w are dense M -coordinate and N -coordinate vectors, respectively. We denote by n_{NZ} the number of nonzero elements in A . Note that, in our case, no specific assumption is made about the structure of the matrix (such as the locations of the non-zeros).

We compute the dot product between each row of A and the vector v . The basic operation is of the form $X \leftarrow (X + \lambda Y) \bmod \ell$, where λ is a non-zero of A , and X and Y are coordinates of the vectors w and v , respectively. To minimize the number of costly reductions modulo ℓ , we accumulate computations, and postpone the final modular reduction of the result. When iterating many products (computations of the form $A^i v$), we can further accumulate several SpMVs before reducing, as long as the intermediate results do not exceed the largest representable integer. As far as arithmetic over $\mathbb{Z}/\ell\mathbb{Z}$

is concerned, we chose to use the Residue Number System, which appears to be more suited to the fine grained parallelism inherent to the SPMT computing model than the usual multi-precision representation of large integers. A comparison of the two representations will be proposed in subsection V-D.

III. RESIDUE NUMBER SYSTEM AND MODULAR ARITHMETIC

A. A brief recall

The Residue Number System (RNS) is based on the Chinese Remainder Theorem (CRT). Let $\mathcal{B} = (p_1, p_2, \dots, p_n)$ be a set of co-prime integers, which we call an *RNS-basis*. We define P as the product of all the p_i 's. The RNS uses the fact that an integer X within $[0, P - 1]$ can be uniquely represented by the list (x_1, x_2, \dots, x_n) , where x_i is the residue of X modulo p_i , which we write as $x_i = |X|_{p_i}$.

This number system is particularly interesting for arithmetic over large integers, since it distributes the computation over several small residues. In other words, the computational units that will work on the residues are independent and need no synchronization nor communication, as there is no carry propagation [16], [17].

If X and Y are given in their RNS representations $X_{RNS} = (x_1, \dots, x_n)$ and $Y_{RNS} = (y_1, \dots, y_n)$, according to \mathcal{B} , and such that $X, Y < P$, RNS addition and multiplication are realized by modular addition and multiplication on each component:

$$X_{RNS} +_{RNS} Y_{RNS} = (|x_1 + y_1|_{p_1}, \dots, |x_n + y_n|_{p_n})$$

$$X_{RNS} \times_{RNS} Y_{RNS} = (|x_1 \times y_1|_{p_1}, \dots, |x_n \times y_n|_{p_n})$$

The result (e.g., $X + Y$) should belong to the interval $[0, P - 1]$ if we want to obtain a valid RNS representation. Otherwise, it will be reduced modulo P .

B. RNS on GPU

We represent the finite ring $\mathbb{Z}/\ell\mathbb{Z}$ as the integer interval $[0, \ell - 1]$. Each element is stored in its RNS form. For performance considerations, we opted for 64-bit moduli. Each RNS residue is represented as an unsigned 64-bit integer in the range $[0, p_i - 1]$.

We implement the RNS operations in PTX (*parallel thread execution*) pseudo-assembly language for CUDA [18]. The basic RNS operation is $z_i \leftarrow (x_i + \lambda y_i) \bmod p_i$, where $0 \leq x_i, y_i, z_i < p_i$ are RNS residues. To speed up the reduction modulo p_i , the moduli are of the pseudo-Mersenne form $2^{64} - c_i$, with c_i as small as possible. With this particular form, the reduction can be done in a small number of additions and (half) products [19].

To make our implementation more efficient, we propose to relax the condition on the input and output: $x_i, z_i \in [0, 2^{64} - 1]$. Since $2^{64} \equiv c_i \pmod{p_i}$, we consider the high part of $x_i + \lambda y_i$ — i.e the part that exceeds 2^{64} —, which we denote by t . We have $(x_i + \lambda y_i) \equiv ((x_i + \lambda y_i) \bmod 2^{64}) + t \times c_i \pmod{p_i}$. If the right-hand side exceeds 2^{64} , as it is below than 2×2^{64} , c_i is in the same way added to the low part. The comparison

with p_i is thus replaced by the overflow detection, which is cheaper, and only c_i is stored instead of p_i (cf. Algo. 2).

For the problem sizes considered, taking $c_i < 2^8$ is sufficient. The algorithm is valid while $|t \times c_i| < 2^{64}$. This ensures correctness for $|\lambda|$ up to 2^{56} .

Algorithm 1: RNS Add

Input : $0 \leq x_i < 2^{64}, 0 \leq y_i < p_i$
Input : $c_i < 2^8$, such that $p_i = 2^{64} - c_i$
Output: $z_i \equiv x_i + y_i \pmod{p_i}, 0 \leq z_i < 2^{64}$
1 $z_i \leftarrow |x_i + y_i|_{2^{64}}$ // $x_i + y_i < 2^{64} + p_i$
2 **if** carry **then**
3 $z_i \leftarrow z_i + c_i$

Algorithm 2: RNS AddMul

Input : $0 \leq x_i < 2^{64}, 0 \leq y_i < p_i, 0 \leq \lambda < 2^{56}$
Input : $c_i < 2^8$, such that $p_i = 2^{64} - c_i$
Output: $z_i \equiv x_i + \lambda \times y_i \pmod{p_i}, 0 \leq z_i < 2^{64}$
1 $z_i \leftarrow (x_i + \lambda \times y_i)_L$ // the low 64-bit part
2 $t \leftarrow (x_i + \lambda \times y_i)_H$ // the high 64-bit part
3 $z_i \leftarrow |z_i + c_i \times t|_{2^{64}}$ // $c_i \times t < 2^{64}$
4 **if** carry **then**
5 $z_i \leftarrow z_i + c_i$

C. RNS modular reduction

In the chosen RNS representation, $(P - 1)$ is the largest representable integer. So in the case of iterative SpMV, we can accumulate at most $\log(\frac{P}{\ell-1}) / \log(r)$ matrix-vector products before having to reduce modulo ℓ , where r corresponds to the largest row norm (defined as the sum of the absolute values of its elements) in the matrix. To reduce the vector dst modulo ℓ , there are two options:

- Convert it to the conventional form (*reconstruction*), then reduce it and convert the reduced vector back to its RNS form (*evaluation*).
- Perform the modular reduction in RNS.

For the first option, the steps of reconstruction, reduction and evaluation are time-consuming (much longer than the SpMV). To further increase the number of possible products before reduction, we can increase P , but this would make the RNS operations more expensive.

The second option allows us to remain in the RNS form throughout the solver computation. However, the RNS representation has the drawback that modular reduction is difficult to perform, compared to addition and multiplication. There are in the literature algorithms for RNS modular multiplication [20], [21], that we adapt here to derive a method to perform modular reduction.

We start from the CRT reconstruction:

$$X = \sum_{i=1}^n \gamma_i P_i \pmod{P},$$

where we have defined $P_i = P/p_i$ and $\gamma_i = |x_i P_i^{-1}|_{p_i}$.

Let us also define the integer α as follows

$$\alpha = \left\lfloor \sum_{i=1}^n \frac{\gamma_i P_i}{P} \right\rfloor = \left\lfloor \sum_{i=1}^n \frac{\gamma_i}{p_i} \right\rfloor. \quad (1)$$

X can be written as $\sum_{i=1}^n \gamma_i P_i - \alpha P$ and, since $\gamma_i < p_i$, we have that $0 \leq \alpha < 1$.

Since $p_i \approx 2^k$, in [22], Kawamura et al. approximate the quotient γ_i/p_i using only the s most significant bits of $\gamma_i/2^k$. Hence, an estimate for α is proposed as

$$\hat{\alpha} = \left\lfloor \sum_{i=1}^n \frac{\lfloor \frac{\gamma_i}{2^{k-s}} \rfloor}{2^s} + \Delta \right\rfloor, \quad (2)$$

where s is an integer parameter in $[1, k]$ and Δ an error correcting term in $]0, 1[$.

The following result is derived from [22].

Proposition 1. *If $0 \leq X < (1 - \Delta)P$ and $(\epsilon + \delta) \leq \Delta < 1$ where $\epsilon = \sum_{i=1}^n \frac{c_i}{2^k}$ and $\delta = n \frac{2^{k-s} - 1}{2^k}$, then $\alpha = \hat{\alpha}$.*

Proof: To measure the inaccuracy induced by the approximation of γ_i/p_i , we define the error terms

$$\epsilon_i = \frac{\gamma_i}{p_i} - \frac{\gamma_i}{2^k} \text{ and } \delta_i = \frac{\gamma_i}{2^k} - \frac{\lfloor \frac{\gamma_i}{2^{k-s}} \rfloor}{2^s}.$$

We have

$$\sum_{i=1}^n \frac{\gamma_i}{p_i} - \left(\sum_{i=1}^n \epsilon_i + \sum_{i=1}^n \delta_i \right) = \sum_{i=1}^n \frac{\lfloor \frac{\gamma_i}{2^{k-s}} \rfloor}{2^s}. \quad (3)$$

As $\epsilon_i, \delta_i \geq 0$,

$$\sum_{i=1}^n \frac{\lfloor \frac{\gamma_i}{2^{k-s}} \rfloor}{2^s} \leq \sum_{i=1}^n \frac{\gamma_i}{p_i}. \quad (4)$$

Since $\gamma_i < p_i$, an upper bound is given by

$$\epsilon_i = \frac{\gamma_i}{p_i} \left(\frac{2^k - p_i}{2^k} \right) < \frac{2^k - p_i}{2^k} = \frac{c_i}{2^k}.$$

Moreover, as $\gamma_i - \lfloor \frac{\gamma_i}{2^{k-s}} \rfloor \times 2^{k-s} \leq 2^{k-s} - 1$, we obtain

$$\delta_i = \frac{\gamma_i - \lfloor \frac{\gamma_i}{2^{k-s}} \rfloor \times 2^{k-s}}{2^k} \leq \frac{2^{k-s} - 1}{2^k}.$$

Eq. 4 and the introduction of the previous inequalities in Eq. 3 lead to

$$\sum_{i=1}^n \frac{\gamma_i}{p_i} - (\epsilon + \delta) < \sum_{i=1}^n \frac{\lfloor \frac{\gamma_i}{2^{k-s}} \rfloor}{2^s} \leq \sum_{i=1}^n \frac{\gamma_i}{p_i}.$$

The substitution of $\alpha + \frac{X}{P}$ for $\sum_{i=1}^n \frac{\gamma_i}{p_i}$ and the addition of Δ result in:

$$\alpha + \frac{X}{P} - (\epsilon + \delta) + \Delta < \sum_{i=1}^n \left\lfloor \frac{\gamma_i}{2^{k-s}} \right\rfloor + \Delta \leq \alpha + \frac{X}{P} + \Delta. \quad (5)$$

As $X < (1 - \Delta)P$, it follows that:

$$\alpha < \sum_{i=1}^n \left\lfloor \frac{\gamma_i}{2^s} \right\rfloor + \Delta < \alpha + 1. \quad \blacksquare$$

Now, let us define Z as $\sum_{i=1}^n \gamma_i |P_i|_\ell - |\alpha P|_\ell$. We can easily check that Z lives in $[0, n2^{k\ell} - 1]$ and is congruent to X modulo ℓ . After determining α , we are able to perform a fully RNS computation of Z . Algorithm 3 summarizes the steps of computation. We denote by \hat{P} and Γ the RNS forms ($|P_1^{-1}|_{p_1}, \dots, |P_n^{-1}|_{p_n}$) and $(\gamma_1, \dots, \gamma_n)$, respectively.

Algorithm 3: RNS modular reduction

Input : $k, \ell, n, \Delta, s, \mathcal{B}, \hat{P}$ precomputed
Input : precomputed table $|P_i|_\ell$ for $i = 1 \dots n$
Input : precomputed table $|\alpha P|_\ell$ for $\alpha = 1 \dots n - 1$
Input : positive $X < (1 - \Delta)P$ in RNS form
Output: $Z \equiv X \pmod{\ell}, Z < n2^{k\ell}$

- 1 $\Gamma \leftarrow X \times \hat{P}$ //1 RNS multiplication
- 2 **for** $i = 1 \dots n$ **do**
- 3 $C_i \leftarrow \gamma_i \times |P_i|_\ell$ //1 RNS multiplication
- 4 $Z \leftarrow \sum_{i=1}^n C_i$ // $(n - 1)$ RNS additions
- 5 $\alpha \leftarrow \left\lfloor \sum_{i=1}^n \frac{\gamma_i}{2^{k-s}} \right\rfloor + \Delta$ //sum of n s -bit terms
- 6 $Z \leftarrow Z - |\alpha P|_\ell$ //1 RNS subtraction

All the operations can be evaluated in parallel on the residues, except for step 3, where a broadcast of the whole vector Γ is needed.

Even if the obtained result Z is not the exact reduction of X , it is bounded by $n2^{k\ell}$. So, we guarantee that the intermediate results of the SpMV computation do not exceed a certain bound less than P . Notice that this RNS reduction algorithm imposes that P be one modulus (k bits) larger than implied by the earlier condition $\ell < P$.

IV. SPARSE MATRIX STORAGE FORMATS

A. Compressed Sparse Row (CSR)

The CSR format stores the column indices and the values into two arrays of n_{NZ} elements: `id` and `data`. A third array of row pointers, `ptr`, of length $N + 1$, is used to indicate the beginning and the end of each row. Non-zeros are sorted by their row indices. The CSR format eliminates the explicit storage of the row index, and is convenient for a direct access to the matrix, since `ptr` indicates where each row starts and

ends in the other two ordered arrays. In the pseudo-code, the vectors `src` and `dst` represent v and w , respectively.

1) *Scalar*: To parallelize the product for the CSR format, a simple way is to assign one thread for each row (*scalar* approach) (cf. listing 1). Temporary results are stored on registers and the final result is written to global memory. In the running sum, the RNS AddMul is performed in PTX. For simplicity, it is denoted here by the operators “*”, “+” and “%”.

The major drawback of the scalar approach is that making each thread iterate over all the RNS residues increases the number of registers per thread. This approach also suffers from a poor global memory load efficiency, because the threads within a same warp access in a non-contiguous fashion `id` and `data`.

```

__global__ void
spmv_csr_scalar_kernel( const int * data, const int * id,
                       const int * ptr, const int * p,
                       const uint2 * src, uint2 * dst )
{
    // one thread per row
    int row = BLOCK_SIZE * blockIdx.x + threadIdx.x;

    if ( row >= nRows ) return;
    uint2 vals[n]; // array of n (the # of RNS moduli) 64-bit words

    for( int j = 0; j < n; ++j )
        vals[j] = 0; // initialization

    // running sum
    for( int i = ptr[row]; i < ptr[row+1]; ++i )
        for( int j = 0; j < n; ++j )
            vals[j] = ( vals[j] + data[i] * src[id[i] * n + j] ) % p[j];

    for( int j = 0; j < n; ++j )
        dst[row * n + j] = vals[j]; // from registers to global memory
}

```

Listing 1. SpMV kernel for the scalar CSR format

2) *Vector*: The *vector* approach consists in assigning a warp to each row of the matrix. The threads within a warp access neighboring non-zeros elements, which makes the warp accesses to `id` and `data` contiguous. Each thread computes its partial result in shared memory, then a parallel reduction in shared memory is required to combine the per-thread results. No synchronization is needed, since threads belonging to a same warp are implicitly synchronized.

However, in the context of RNS arithmetic, the *vector* scheme still suffers from a low load/write efficiency when accessing `src` and `dst`, because threads within the same warp simultaneously access residues of different entries, which are not contiguous.

3) *Residue-vector*: To overcome the limitations of the previous approaches, we propose to organize the threads within a warp into n_{GPS} groups of n threads, where $n_{GPS} \times n$ (that we denote by n_{CUS}) is closest to 32. Each group is associated to a non-zero (cf. Listing 2). For instance for $n = 6$, we take $n_{GPS} = 5$, so the first 6 threads process in parallel residues of the 1st non-zero, threads 6 to 11, process the 2nd non-zero, and so on, and we will have two idle threads. Like for the vector approach, a reduction is needed to combine the results of threads working on the same residue and belonging to different groups.

In this scheme, which we call *residue-vector*, the number of registers is reduced by eliminating the per-thread loop to process all the residues of an entry, and contiguous accesses to `src` and `dst` are performed.

```

__global__ void
spmv_csr_residue_vector_kernel ( const int * data, const int * id,
                                const int * ptr,  const int * p,
                                const uint2 * src,  uint2 * dst )
{
    __shared__ uint2 vals [(BLOCK_SIZE / 32) * NCUS];

    int tid = BLOCK_SIZE * blockIdx.x + threadIdx.x;
    int warp_id = tid / 32;
    int lane = tid & 31;
    int resid_id = lane % n; // residue index within the entry
    int vals_id = (threadIdx.x / 32) * NCUS + lane;

    int row = warp_id; // one warp per row

    if (row >= nRows) return;
    if (lane >= NCUS) return; //idle threads
    uint2 val = 0;

    for (int i = ptr[row] + lane / n; i < ptr[row+1]; i+=NGPS)
        val = (val + data[i] * src[id[i] * n + resid_id]) % p[resid_id];

    vals[vals_id] = val; // from registers to shared memory

    if (lane < n) { // first group
        for (int k = 1; k < NGPS; ++k)
            val = (val + vals[vals_id + k * n]) % p[resid_id]; // reduction
        dst[row * n + residue_id] = val;
    }
}

```

Listing 2. SpMV kernel for the residue-vector CSR format

B. Coordinate (COO)

For the COO format, the row/column indices and the value are explicitly stored to specify a non-zero entry. The format consists of three arrays `row_id`, `col_id` and `data` of n_{NZ} elements. In this work, we propose to sort the entries by their row index.

A typical way to work with the COO format on GPU is to assign one thread to each entry. This implies that different threads from different warps will process the same row. To combine their results, one possibility is to do atomic updates on global memory to the result vector `dst`, which significantly decreases the performances. Another possibility is that each thread computes its partial dot product, then performs a segmented reduction [23], [24] to sum the partial results of the other threads belonging to the same warp and spanning the same row. We followed the scheme proposed by the library CUSP¹, which performs the segmented reduction in shared memory, using the row indices as segment descriptors. As for the *residue-vector* CSR kernel, we assign a group of threads to each non-zero. Each warp iterates over its interval, processing NGPS elements at a time. If a spanned row is terminated, its result is written to `dst`, otherwise, the row index and the partial dot product are stored in temporary arrays. Then, a second kernel performs the combination of the per-warp results.

¹<http://code.google.com/p/cusp-library/>

C. Sliced Coordinate (SLCOO)

The SLCOO format was inspired from the CADO-NFS² software for CPUs. It was introduced for GPUs by Schmidt et al. for integer factorization, in the particular case of matrices over GF(2) [14]. The aim of this format is to increase the cache hit rate that limits the CSR and COO performances. Like COO, the SLCOO representation stores the row indices, column indices and values. It divides the matrix into horizontal slices, where the non-zeros of a slice are sorted according to their column indices with the aim to reduce the irregular accesses on source vector `src`, if they had been sorted by their row indices. A fourth array `ptrSlice` indicates the beginning and end of each slice.

For the SLCOO kernel (cf. Listing 3), one warp works on a slice. A group of n threads processes in parallel a non-zero, each thread being assigned to a residue. Since each thread works on more than one row, it needs to have individual storage for its partial per-row results, or to be able to have exclusive access to a common resource. In [14], where a thread block had been assigned to each slice, three possibilities have been mentioned to solve this issue:

- Small SLCOO: each thread has one exclusive entry in shared memory to store the partial result for each row.
- Medium SLCOO: threads having the same lane (index within the warp) share one entry per row in shared memory and access it by an atomic XOR operation.
- Large SLCOO: all threads share one entry per row.

The Medium SLCOO allows one to put $(BLOCK_SIZE/32)$ times as many rows on one slice than the Small SLCOO. Similarly, the Large SLCOO allows one to fill 32 times as many rows than the Medium SLCOO. This way, increasing the number of rows per slice (from Small to Large variants) increases the texture cache hit rate, which compensates the drawbacks of the atomic accesses.

However, because of the reduction, it is not possible to perform an atomic addition in RNS. For this reason, we only implement the Small SLCOO. Furthermore, this kernel suffers from the excessive usage of shared memory, since it stores $\mathbb{Z}/\ell\mathbb{Z}$ entries.

```

__global__ void
spmv_slcoo_kernel(const int * data, const int * row_id,
                  const int * col_id, const int * ptrSlice,
                  const int * p, const uint2 * src,
                  uint2 * dst)
{
    __shared__ uint2 vals [SLICE LENG * (BLOCK_SIZE / 32) * NCUS];

    int tid = BLOCK_SIZE * blockIdx.x + threadIdx.x;
    int warp_id = tid / 32;
    int lane = tid & 31;
    int resid_id = lane % n;
    int vals_id = (threadIdx.x / 32) * NCUS + lane;

    if (lane >= NGPS * n) return;

    for (int j = 1; j < SLICE LENG; ++j)
        vals[j * (BLOCK_SIZE / 32) * NCUS + vals_id] = 0;

    for (int i = ptrSlice[warp_id] + lane / n; i < ptrSlice[warp_id + 1];
         i+= NGPS)

```

²<http://cado-nfs.gforge.inria.fr/>

```

vals [(row_id[i] % SLICE_LENG) * (BLOCK_SIZE / 32) * NCUS +
      vals_id] += data[i] * src[col_id[i] * n + resid_id] % p[resid_id];

if (lane < n) {
  // reduction in shared memory
  for (int j = 1; j < SLICE_LENG; ++j)
    for (int k = 1; k < NGPS; ++k)
      vals [j * (BLOCK_SIZE / 32) * NCUS + vals_id] += vals[j * (
        BLOCK_SIZE / 32) * NCUS + vals_id + k * n] % p[
        resid_id];
  // from shared memory to global memory
  for (int j = 1; j < SLICE_LENG; ++j)
    dst [(warp_id * SLICE_LENG + j) * n + resid_id] = vals[j * (
      BLOCK_SIZE / 32) * NCUS + vals_id];
}
}

```

Listing 3. SpMV kernel for the SLCOO format

There are in the literature other SpMV formats that we did not consider such as DIA (Diagonal), ELL (ELLPACK), etc. The DIA stores the offset of each sub-diagonal from the main diagonal and is appropriate only for matrices that satisfy some sparsity patterns, which is not our case. The ELL format extends the CSR arrays to N -by- K arrays, where K corresponds to the maximum number of non-zeros per row. For the matrices that we use, this storage scheme provide no advantage over the CSR scheme.

V. COMPARATIVE ANALYSIS OF SPMV KERNELS

All the experiments are run on a system comprising an NVIDIA GeForce GTX 580 and an Intel Xeon Processor E5440 (2.83 GHz). Each SpMV kernel is executed 100 times. We report the computational throughput in terms of GFLOP/s, which we determine by dividing the number of required operations (twice the number of non-zeros multiplied by $2 \times n$) by the running time. The memory throughput is computed by dividing the total number of bytes read/written by all the threads by the running time. Our measurements do not include the time spent to copy data between the host and the GPU, since the matrix and vectors do not need to be transferred back and forth in each SpMV iteration. These delays are thus dwarfed by the computation latencies. The timings do not include the reduction modulo ℓ , since it happens only once every few iterations and also because our implementation based on multi-precision representation does not support it.

The reported measurements are based on the Nvidia Visual Profiler³ and the NVIDIA CUDA Occupancy Calculator⁴ results. Table I summarizes the test matrix over $\mathbb{Z}/\ell\mathbb{Z}$, where ℓ is a prime of up to 280 bits. The matrix was obtained from a DLP resolution program based on the FFS algorithm. The $\mathbb{Z}/\ell\mathbb{Z}$ entries fit in five 64-bit residues. There is an extra 64-bit residue, needed for the modular reduction: $n = 6$.

A. comparison of CSR variants

The scalar kernel uses a large number of registers, which limits the maximum number of warps that can be put on an SM to 24 active warps (per SM), out of the 48 supported. This is reported on the *theoretical occupancy* column in Table II. On

	FFS-280
Size	$1,732,788 \times 1,732,788$
#Non-zeros	86,639,540
Max (row norm)	374
Percentage of ± 1	93.48%
Size of ℓ (bits)	280
Size of M (bits)	384
Size of $n2^k\ell$ (bits)	346

TABLE I
PROPERTIES OF USED MATRICE

the other hand, the vector kernel suffers from shared memory usage (4.5kB for a 96-thread block), which also limits the maximal reachable occupancy to 50%. For these two kernels, the low occupancy significantly decreases the performances. Compared to the vector kernel, the scalar one reaches higher throughput, since it allows an L1-oriented configuration (48k-L1, 16k-shared) of the on-chip memory, while for the vector kernel, a shared-oriented configuration (16k-L1, 48k-shared) is required.

Concerning the global memory access pattern, the column *Global Load/Store Efficiency* gives the ratio of requested memory transactions to the number of transactions performed, which reflects whether accesses are all perfectly coalesced (100% efficiency) or not. For scalar and vector kernels, uncoalesced accesses cause the bandwidth loss and the performance degradation.

The residue-vector (CSR-RV) kernel satisfies better the GPU architectural characteristics. It makes the write accesses coalesced (100% store efficiency). For the loads, we obtain only 45% efficiency, because a warp accesses several non-zeros, which are not necessarily contiguous (sparsity of A).

B. COO/CSR-RV comparison

Due to the segmented reduction, the COO kernel performs more instructions and requires more registers than the CSR kernel. Thread divergence happens more often, because of the several branches that threads belonging to the same warp can take. Memory access pattern for writing is less efficient, compared to the CSR kernel (cf. *Store Efficiency* in Table II), due to the fact that some results are written to *dst*, the rest to the temporary vector. Consequently, the COO kernel reaches lower throughput than the CSR one.

C. SLCOO/CSR-RV comparison

Table III compares several SLCOO kernels for different slice lengths. By making the slices larger, we increase the usage of shared memory proportionally to the slice length. For slice length above 4, the requested shared memory exceeds 16kB, two options are possible:

- Decrease the number of blocks on a multiprocessor and keep the 48k-L1-16k-shared configuration (cf. kernels marked † in Table III). This improves the cache hit rate.
- Keep the same occupancy and switch to 16k-L1-48k-shared configuration. This limits the cache hit rate.

³<http://developer.nvidia.com/nvidia-visual-profiler>⁴<http://developer.nvidia.com/cuda-downloads>

	Registers per Thread	Shared Mem. per SM	Branch Divergence	(Theoretical) Occupancy	Global Load / Store Efficiency	Timing in ms	Global Load / Store Throughputs	Computation. Throughput
CSR-S	37	0	48.2%	(50%) 45.9%	4.2% / 12.3%	187.23	34.95 / 0.42	11.1
CSR-V	28	36,864	59.8%	(50%) 36.4%	7.5% / 25%	216.48	48.79 / 0.36	9.6
CSR-RV	16	11,520	25.9%	(100%) 69.6%	46.9% / 100%	59.36	179.4 / 6.52	35.02
COO	23	13,824	36%	(66.7%) 64.9%	44% / 37.5%	110.47	88.36 / 0.72	18.82

TABLE II
COMPARISON OF *scalar*, *vector* AND *residue-vector* CSR-KERNELS AND COO KERNEL.

For slices of length 2, the 48k-L1-16k-Shared configuration is sufficient, but we can not reach higher occupancy, due to the high number of registers per thread (24).

For slices of length 12, only the 16k-L1-48k-Shared configuration is possible, since each block exceeds the threshold of 16kB.

We remark that for both configurations, increasing the slice length improves the cache hit rate, since accesses on the source vector are less irregular. However, the limitation of the occupancy yields poor performance compared to the CSR-RV kernel.

D. RNS and Multi-precision arithmetics comparison

The idea behind the use of RNS arithmetic rather than multi-precision (MP) arithmetic is that RNS can significantly decrease data sharing between the threads and arithmetic operations required for the carry generation/propagation (cf. Table IV). Consequently, it allows us to reach higher occupancy and better performance.

The low difference between the two kernels (< 5%) is due to the fact that the RNS kernel requires an extra residue for the modular reduction step. If that residue is removed, the speed-up of RNS compared MP is around 15%.

To perform the modular reduction in RNS, we choose $\Delta = 0.25$. For the maximum row norm that we have (374), this allows to do up to 4 iterative SpMV's before having to reduce. The reduction kernel takes 5.7 ms, which corresponds to 1.42 ms per iteration ($\sim 2.5\%$ of the SpMV delay).

VI. IMPROVEMENTS ON CSR KERNEL

To improve the kernel performance, one should take into account the GPU architectural characteristics: the management of the memory accesses and the partitioning of the computations. The effects and the results corresponding to each improvement are reported in Table V.

A. Texture caching

Although our SpMV kernel suffers from irregular load accesses, a thread is likely to read from an address near the addresses that nearby threads (of the same group) read. For this reason, we enable the texture caching by binding the source vector on texture memory and by replacing reads of the form `src[j]` with a texture fetch `tex1Dfetch(src_tex, j)`. This improves the global memory efficiency and consequently the throughput.

B. Reordering the non-zeros of a row

When computing discrete logarithms, most of the coefficients of the matrix are ± 1 . It seems promising to treat multiplications by these coefficients differently from other coefficients: additions and subtractions are less expensive than multiplications. Moreover, we are not able to produce the same code for positive and negative values, so negative coefficients are processed differently as well. All these separations result in code divergence, that we fix by reordering the non-zeros in the matrix such that values of the same subgroup are contiguous. This decreases the branch divergence and increases the total throughput.

C. Improving warp balancing

In the proposed kernel, each warp processes a single row. This requires launching a large number of warps. Consequently, there is a delay to schedule those launched warps. Instead, we propose that each warp iterates over a certain number of rows. Although this increases the usage of registers (18 \rightarrow 24), it improves the achieved occupancy. To further increase the occupancy, we permute the rows such that each warp roughly gets the same work load.

VII. CONCLUSION

We have investigated different data structures to perform iterative SpMV for DLP matrices on GPUs. We have adapted the kernels for the context of large finite fields and added optimizations suitable to the sparsity and the specific computing model. The CSR format based on the *residue-vector* approach appears to be the most efficient one. The SLCOO poses for the sizes that we use some hardware difficulties that nullify its contribution on increasing the cache hit rate. Future GPUs may enhance the performances. We have shown that using RNS for finite field arithmetic provides a considerable degree of independence, which can be exploited by massively parallel hardware.

REFERENCES

- [1] A. M. Odlyzko and A. M. Odlyzko, "Discrete logarithms in finite fields and their cryptographic significance," 1984.
- [2] J. M. Pollard, "A monte carlo method for factorization," *BIT Numerical Mathematics*, vol. 15, pp. 331–334, 1975.
- [3] D. Shanks, "Class number, a theory of factorization, and genera," in *1969 Number Theory Institute (Proc. Sympos. Pure Math., Vol. XX, State Univ. New York, Stony Brook, N.Y., 1969)*, Providence, R.I., 1971, pp. 415–440.

	Slice Length	L1 Cache Hit Rate	Shared Memory per Block	Blocks per SM	(Theoretical) Occupancy	Timing in ms	Global Load / Store Throughputs	Computational Throughput
SLCOO	2†	41%	3840	4	(66.7%) 57.3%	64.45	153.63 / 1.2	32.26
	4†	45.4%	7680	2	(33.3%) 29.8%	79.91	122.53 / 0.99	26.02
	4	15.7%	7680	4	(66.7%) 61.1%	71.77	136.43 / 1.08	28.97
	8†	46.3%	15360	1	(16.7%) 16.1%	130.23	74.74 / 0.61	15.96
	8	19%	15360	3	(50%) 46.8%	73.2	132.96 / 1.06	28.4
	12	26.9%	23040	2	(33.3%) 31.3%	85.63	113.48 / 0.92	24.28
CSR-RV	-	48.8%	1920	6	(100%) 69.6%	59.36	179.4 / 6.52	35.02

TABLE III
COMPARISON OF CSR-RV KERNEL AND SEVERAL SLCOO KERNELS FOR DIFFERENT SLICE LENGTHS.

	Registers per Thread	Shared Memory per Block	Executed Instructions	(Theoretical) Occupancy	Timing in ms	Computational Throughput
MP	20	2880	$8.05 \cdot 10^8$	(83.3%) 47.4%	62.07	33.5
RNS	16	1920	$7.73 \cdot 10^8$	(100%) 69.6%	59.36	35.02

TABLE IV
COMPARISON OF RNS AND MP KERNELS FOR THE CSR-RV FORMAT.

	Performance Effects	Timing in ms	Throughput (speedup)
Texture caching	Global Load Efficiency: 46.9% \rightarrow 84.7%	44.33	46.9 (+33%)
Non-zeros reordering	Branch Divergence: 25.9% \rightarrow 12.9%	41.66	49.91 (+6.4%)
Multiple iterations	(Theoretical) Occupancy: (100%) 69.6% \rightarrow (83.3%) 74.8%	41.5	50.1 (+0.4%)
Rows permutation	(Theoretical) Occupancy: (83.3%) 74.8% \rightarrow (83.3%) 82.5%	41.04	50.66 (+1.1%)

TABLE V
PERFORMANCE EFFECTS OF THE IMPROVEMENTS ON THE CSR-RV KERNEL AND THEIR SPEEDUPS.

- [4] L. Adleman, "A subexponential algorithm for the discrete logarithm problem with applications to cryptography," in *Proceedings of the 20th Annual Symposium on Foundations of Computer Science*. Washington, DC, USA: IEEE Computer Society, 1979, pp. 55–60.
- [5] B. A. LaMacchia and A. M. Odlyzko, "Solving large sparse linear systems over finite fields," in *CRYPTO '90*. London, UK: Springer-Verlag, 1991, pp. 109–133.
- [6] T. Hayashi, T. Shimoyama, N. Shinohara, and T. Takagi, "Breaking pairing-based cryptosystems using η_t pairing over $GF(3^{97})$," *Cryptology ePrint Archive*, Report 2012/345, 2012.
- [7] C. Lanczos, "Solution of systems of linear equations by minimized iterations," *J. Res. Natl. Bur. Stand.*, vol. 49, pp. 33–53, 1952.
- [8] D. H. Wiedemann, "Solving sparse linear equations over finite fields," *IEEE Trans. Inf. Theor.*, vol. 32, no. 1, pp. 54–62, Jan. 1986.
- [9] C. Pomerance and J. W. Smith, "Reduction of huge, sparse matrices over finite fields via created catastrophes," *Experiment. Math.*, vol. 1, pp. 89–94, 1992.
- [10] *CUDA Programming Guide Version 4.2*, 2012, <http://developer.nvidia.com/cuda-downloads>.
- [11] N. Bell and M. Garland, "Efficient sparse matrix-vector multiplication on CUDA," NVIDIA Corporation, Technical Report NVR-2008-004, Dec. 2008.
- [12] —, "Implementing sparse matrix-vector multiplication on throughput-oriented processors," in *SC '09: Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*. New York, NY, USA: ACM, 2009, pp. 1–11.
- [13] F. Vázquez, E. M. Garzón, J. A. Martínez, and J. J. Fernández, "The sparse matrix vector product on GPUs," University of Almeria, Technical report, Jun. 2009.
- [14] B. Schmidt, H. Aribowo, and H.-V. Dang, "Iterative sparse matrix-vector multiplication for integer factorization on GPUs," in *Euro-Par 2011 Parallel Processing*, vol. 6853. Springer, 2011, pp. 413–424.
- [15] B. Boyer, J.-G. Dumas, and P. Giorgi, "Exact sparse matrix-vector multiplication on GPU's and multicore architectures," *CoRR*, vol. abs/1004.3719, 2010.
- [16] N. S. Szabo and R. I. Tanaka, *Residue Arithmetic and Its Applications to Computer Technology*. New York: McGraw-Hill Book Company, 1967.
- [17] F. J. Taylor, "Residue arithmetic a tutorial with examples," *Computer*, vol. 17, pp. 50–62, May 1984.
- [18] *PTX: Parallel Thread Execution ISA Version 3.0*, 2012, <http://developer.nvidia.com/cuda-downloads>.
- [19] J. A. Solinas, "Generalized mersenne numbers," CACR Technical Report CORR 99-39, 1999.
- [20] B. Phillips, Y. Kong, and Z. Lim, "Highly parallel modular multiplication in the residue number system using sum of residues reduction," *Applicable Algebra in Engineering, Communication and Computing*, vol. 21, pp. 249–255, 2010.
- [21] P. P. Shenoy and R. Kumaresan, "Fast base extension using a redundant modulus in RNS," *IEEE Trans. Comput.*, vol. 38, no. 2, pp. 292–297, Feb. 1989.
- [22] S. Kawamura, M. Koike, F. Sano, and A. Shimbo, "Cox-rower architecture for fast parallel montgomery multiplication," in *EUROCRYPT'00*. Berlin, Heidelberg: Springer-Verlag, 2000, pp. 523–538.
- [23] G. E. Blelloch, M. A. Heroux, and M. Zaghera, "Segmented operations for sparse matrix computation on vector multiprocessors," School of Computer Science, Carnegie Mellon University, Tech. Rep. CMU-CS-93-173, Aug. 1993.
- [24] S. Sengupta, M. Harris, Y. Zhang, and J. D. Owens, "Scan primitives for GPU computing," in *Graphics Hardware 2007*. ACM, Aug. 2007, pp. 97–106.